Introduction and Overview for FE-D Design Review

K. Einsweiler, LBNL

Scope and Goals of this Review

Overview of system specifications and design of module

Overview of FE-D specifications and design at the top level

Brief Summary of Problems with FE-D1

Scope and Goals of this Review

FE-D1 reticle included many die (10 in total):

- Two pixel FE chips (FE-D). For the FE-D1 run, they were identical. For the FE-D1b run, they were very slightly different (changes in M1 and M2 masks).
- Prototype MCC chip. A prototype of several key elements of final MCC, about 20mm^2 core size. Includes FIFO block for final chip, plus large synthesized command decoder block. Presently have tested 14 die, of which 11 work. Appear to be no problems with this design, including operation at XCK = 80-90MHz.
- Prototype CMOS opto-link chips (one DORIC-p and three VDC-p). VDC-p seems to work well up to about 150MHz. DORIC-p has several problems which are now largely understood, and thought to be related to the design.
- Additional test chips: LVDS buffer for rad-hard test board, PM bar with W/L arrays and special pixel transistors, Analog Test chip with all critical FE-D analog elements. All work well, and transistor parameter measurements suggest the run is slightly faster than typical. Many detailed characterizations of Analog Test Chip.
- Have irradiated several Analog Test chips, and several PM bars in April PS run, and have irradiated several VDC chips to relevant doses in May in PS. Many preliminary on-line results, showing significant problems for doses above 10MRad. Further testing of chips after cool-down is required.

We submitted FE-D1 in early August 99:

- Wafers were received back in late Oct. 99. After preliminary evaluations, we returned half of the wafers to TEMIC.
- Based on early testing, it was clear that most of the circuitry worked. There were some serious problems with yield in the control and readout portions of the chip.
- We believe that, based on measurements and simulations, we have addressed all of our errors in the initial chip design. Most involved missing or mis-sized buffers. We will present this information and analysis today.
- A large effort over the last few months has been made to understand the yield problems, and in particular to see if we could find a technological explanation (something special in our layout that interacts badly with some process features). We will summarize the status of this analysis today as well.

We have prepared a new FE-D2 submission:

- The reticle will contain two versions of FE-D2. The first one (FE-D2D) uses the same design as FE-D1 (quasi-static pixel shift register, dynamic logic in pixel hit logic block). The second one (FE-D2S) is modified. Those circuits that were associated with serious yield problems are replaced by static versions, and other circuitry is removed to make space (in particular, three TDAC bits are dropped).

This Review:

- You will see that there are still open questions about whether the yield of a new FE-D2 submission will be significantly different than that of the original FE-D1 submission.
- Due to the difficulty of working with the FE-D1 chips, we have not been able to pursue the types of testing and evaluation which are needed to evaluate FE-D as a pre-production quality chip. This means in particular, constructing modules with the chips, and irradiating complete chips to design doses.
- It is critical for us to see whether it is possible to fabricate our pixel design with decent yield in DMILL as soon as possible.
- For these reasons, we have proposed a “mini-review”, which does not attempt to do a complete review of everything in the FE-D design. Instead, we propose to present what we know about the FE-D1 chip, and what we propose to do next.
- We welcome all of your wisdom and experience in helping us to get the most possible out of the FE-D2 submission.

System Design of Pixel Module

First requirements analysis and discussions led to module system design in 96:

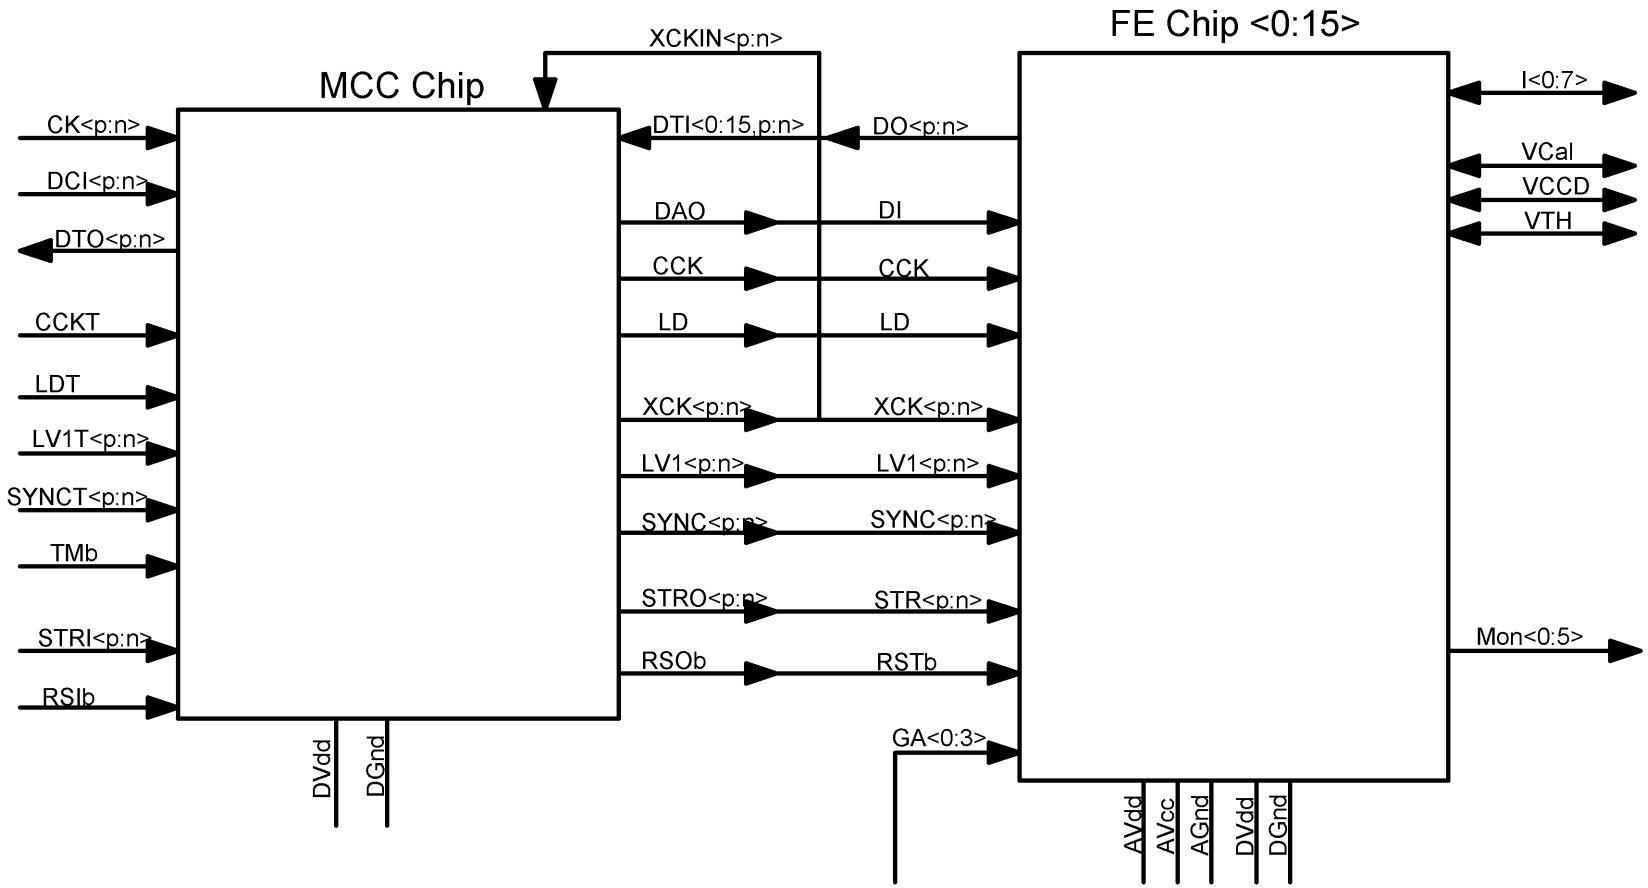

- Two chip design, including a single controller and event-builder chip (MCC), and 16 front-end chips bumped to a single silicon substrate.

Features:

- Basic interface to the outside uses a 3-wire protocol (SerialIn, SerialOut, XCK), which maps onto the SCT opto-link protocol
- Provide a “transparent mode” interface through MCC to FE chips for simplified testing.
- Basic interconnections between FE and MCC use bussed signals. Slow control will not operate when recording events, so it uses full-swing CMOS. All fast signals use low-swing differential “LVDS-like” signaling. Point-to-point signals use 0.5 mA drivers (FE chips only), external or bussed signals use 3 mA drivers.
- To provide enhanced speed and robust module design, the serial output lines are connected from the FE to the MCC in a star topology (16 parallel inputs on MCC).
- There are no remaining analog signals between MCC and FE at this time. All FE chips have internal current references and adjustment DACs to control analog operating points, as well as calibration.
- Architecture is “data-push” style: each crossing for which LV1 accept is present causes all FE chips to autonomously transmit back hit information for the given crossing. LV1 signal may remain set for many contiguous crossings to allow readout of longer time intervals. MCC merges such events together.
- Synchronization signal available to ensure FE chips label LV1 properly.
- System uses two analog supplies to reduce power dissipation.

Requirements Summary

Power budget:

- Nominal total for module is based on 0.6W/cm^2 power density, including all power dissipation in services within the pixel tracker volume. This includes MCC and FE chips, sensor leakage, opto-link, and power transmission.
- A more complete, bottoms-up analysis defined nominal and worst case supply current and voltage for each of three supplies on the FE chip. The typical (worst case) values for the analog supplies are: 10μA at 1.5V (15μA at 1.75V) per pixel for VCCA, and 8μA at 3.0V (12μA at 3.5V) per pixel for VDDA. For the digital supply, allocation for complete chip is total of 25mA at 3.0V (40mA at 4.0V).

Geometry:

- The active die area for the FE chip is 7.2 x 10.8 mm, of which 7.2 x 8.0mm is sensitive area for particle detection. The sensitive area of the FE chips must extend to the edge of the die along 3 sides, with all additional logic and I/O concentrated on the remaining side.
- Physics studies indicate that the pixels should be as narrow as possible in one dimension, and a 50μ pitch has been chosen as reasonably achievable. In the long direction, adequate resolution is obtained with a dimension of 300μ - 400μ. The present prototyping program has frozen the length at 400μ.

Demonstrator Program

- In 1997, we agreed on overall design specifications for the FE and MCC chips necessary to implement this module design in a prototype form. We decided to pursue two prototypes for the FE chip. This was based partly on history and partly on the goal of submitting designs to two rad-hard vendors.
- One was called FE-A, and was designed for submission to AMS 0.8μ BiCMOS. This process was viewed as a prototype vehicle for DMILL. The chip was submitted in Oct. 97, and testing began in Jan. 98. A second, purely CMOS version, FE-C, was submitted in May 98. This chip has 880K transistors.
- The second was referred to as FE-B, and was designed for submission to HP 0.8μ CMOS. This process was viewed as a prototype vehicle for Honeywell. This chip was submitted in Feb. 98, and testing began in Apr. 98. This chip has about 850K transistors.
- A DMILL prototype matrix (no EOC, simple readout) called MAREBO was also submitted in Jul 97 and tested in Jan 98, to verify the FE design in DMILL.
- The MCC was submitted in May 98, along with the FE-C, and returned in late summer. FE-C chips were tested and bonded for evaluation in Sept. testbeams.
- All of these chips contain minor errors, but in all cases their functionality was quite close to the submission goals. Extensive lab testing and testbeam studies have been carried out on all chips. Excellent performance has been achieved.

Demonstrator FE Chip Geometry

- Agreement on pixel size was struck in Sept 96, in order to allow compatible, parallel detector and electronics development.
- The geometry adopted was 50μ x 400μ for the pixel size, with pixels arranged in 18 columns of 160 pixels per column.
- The geometry was mirrored between columns, so that the inputs for pixels in column 0 and 17 are on the outside, and all other columns are paired.
- The input pad geometry in the inner column pairs is then a double row of 50μ pitch pads. The metal pad is specified to be 20μ square, with a 12μ opening in the passivation for the bump-bonding.
- The cut die size must not extend beyond 100μ from the edge of the active area on three sides of the die. Hence, nothing outside of the pixel circuitry is allowed on three sides of the chip, to allow module construction.
- The bottom of the chip (all peripheral logic and I/O pads) are allowed 2800μ, making the total active die region 7.2mm x 10.8mm.
- An I/O pad structure of 48 pads, each consisting of a 100μ x 100μ wire-bond pad, and a group of 4 bump-bond pads, was frozen.

Overview of FE-D

Common design effort, based on combining different aspects of FE-A/B/C chips:

- Front-end design is basically that of FE-C and MAREBO, but is purely CMOS. Although there is a BiCMOS version of this FE (2 bipolars used for discriminator input) which has slightly better timewalk performance than the purely CMOS version, this cell is significantly larger.
- The FE uses a DC-feedback preamp design which provides excellent leakage current tolerance, close to constant-current return to baseline for TOT, and very stable operation with different shaping times.
- It is followed by an AC-coupling stage, and a fast discriminator. A divider network provides fine input baseline (threshold) control and provides the resistance for the high pass filter AC-coupling stage.
- The control logic provides a 3-bit threshold trim capability in each pixel, plus individual mask and calibration inject control. A global FastOR net is created using all pixels enabled for readout, and provides a self-trigger capability.
- All critical bias currents and voltages on the chip are controlled by internal DACs (6 current mode, 2 voltage mode). The current DACs are referenced to an internal V_{BE} reference, and all DACs are controlled via the command decoder.

- Readout design is an upgraded version of the readout architecture of FE-B.
- It uses a 7-bit Grey-coded 40 MHz “timestamp” bus as a timing reference throughout the active matrix. All pixels measure their leading and trailing edge timing by asynchronously latching this reference in RAMs.
- Hits (address plus LE/TE timing) are transferred from the pixels as soon as the trailing edge occurs, using a shared bus structure in the pixel column pair. This bus should operate at transfer rates up to 20 MHz in order to meet our requirements. Low swing signal transmission and sense amplifiers are used to achieve this.
- Significant buffering is provided in the end of column region for hit storage during the L1 latency (up to 3.2μs in ATLAS). Twenty-four buffers are available for each column pair. The coincidence with the L1 trigger is performed in this buffer. Hits from rejected crossings are immediately cleared.
- A readout sequencer stores information on up to 16 events pending readout. As soon as the output serial link is empty, transmission of a new pending event begins. Essentially, sending a L1 trigger corresponds to making a request for the all hits associated with the corresponding beam crossing, which are then pushed off the FE chip to the MCC.

- Global control of the chip is implemented using a simple command protocol. A global register controls Latency, DAC values, enabled columns, clock speeds, and several other parameters.
- A pixel register which snakes through the active array provides access to the 5 control lines in the pixel (Select, Mask, TDAC<0:2>).
- Each chip on a module is geographically addressed, and its identity is controlled by external wire-bonds to avoid confusion.

FE-D Block Diagram:

Basic FE block diagram, expanded from module diagram:

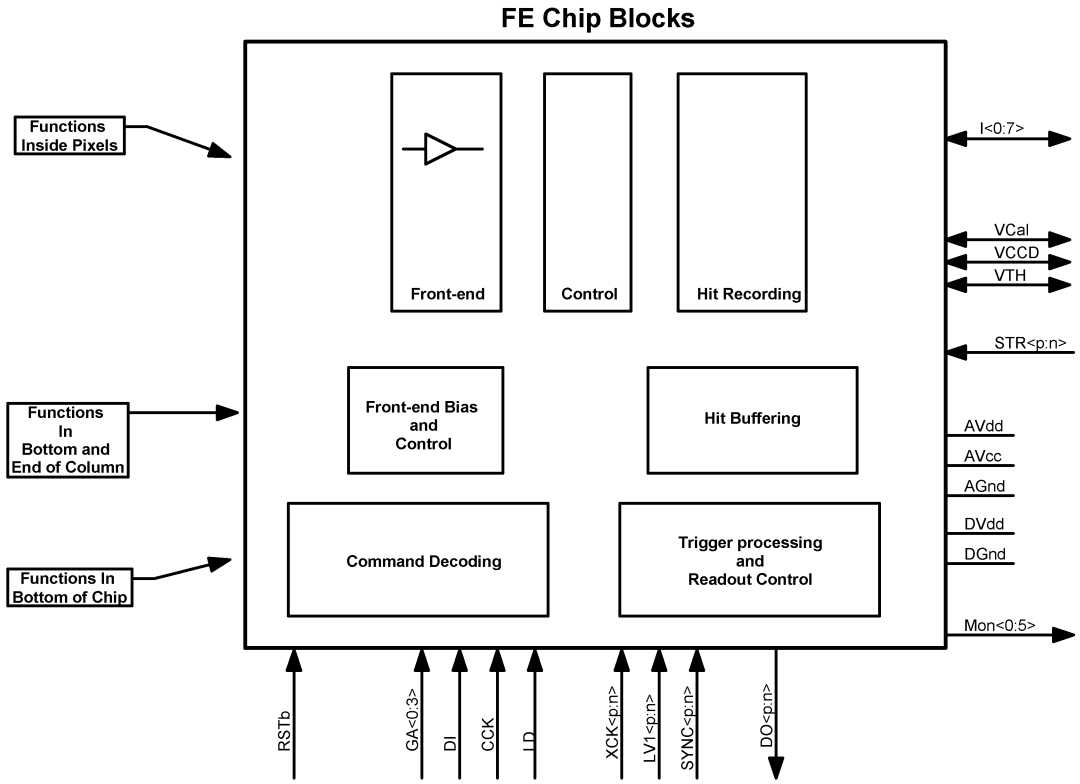

- Within the pixel, there is the front-end (preamp/ discriminator), the control block, and the readout block.
- Just below the active pixel matrix is the biassing and control for the front-end blocks, and the buffering for the readout blocks
- Finally, there is the overall readout control and the command decoding.

- Basic Digital I/O shown on bottom: 4 CMOS inputs for control (RSTb, DI, CCK, LD), and 4 fast, differential I/O's for timing and readout (XCK, LV1, SYNC, DO).
- Calibration and monitoring are shown on the right. A fast, differential strobe (STR) supplies calibration timing, and an analog voltage input (VCal) supplies calibration input. Dedicated monitoring pins include FastOR and test pixel output.

FE-D Pinout and Geometry

Sketch of pin assignments and overall geometry of die:

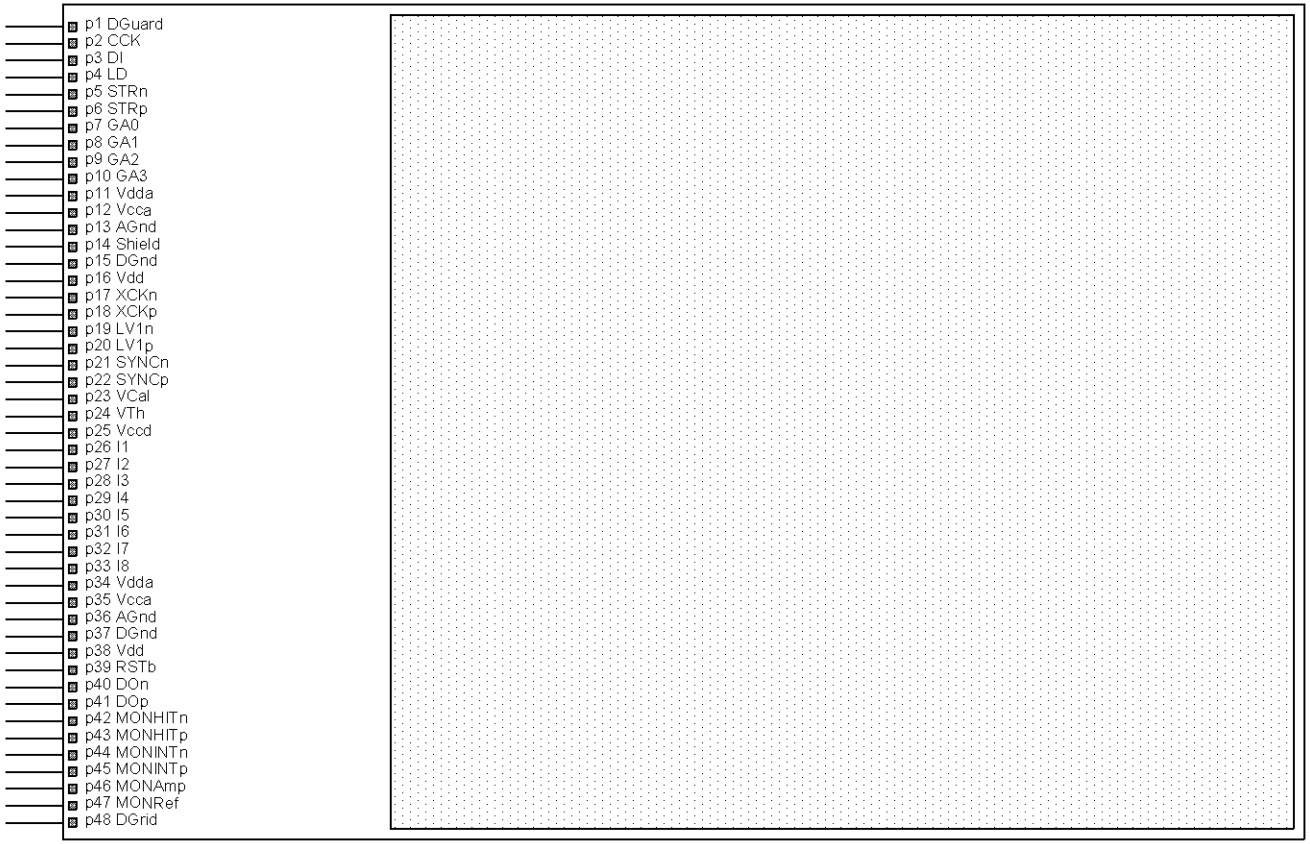

Photo of FE-D1 Mounted on Single Chip Support Card:

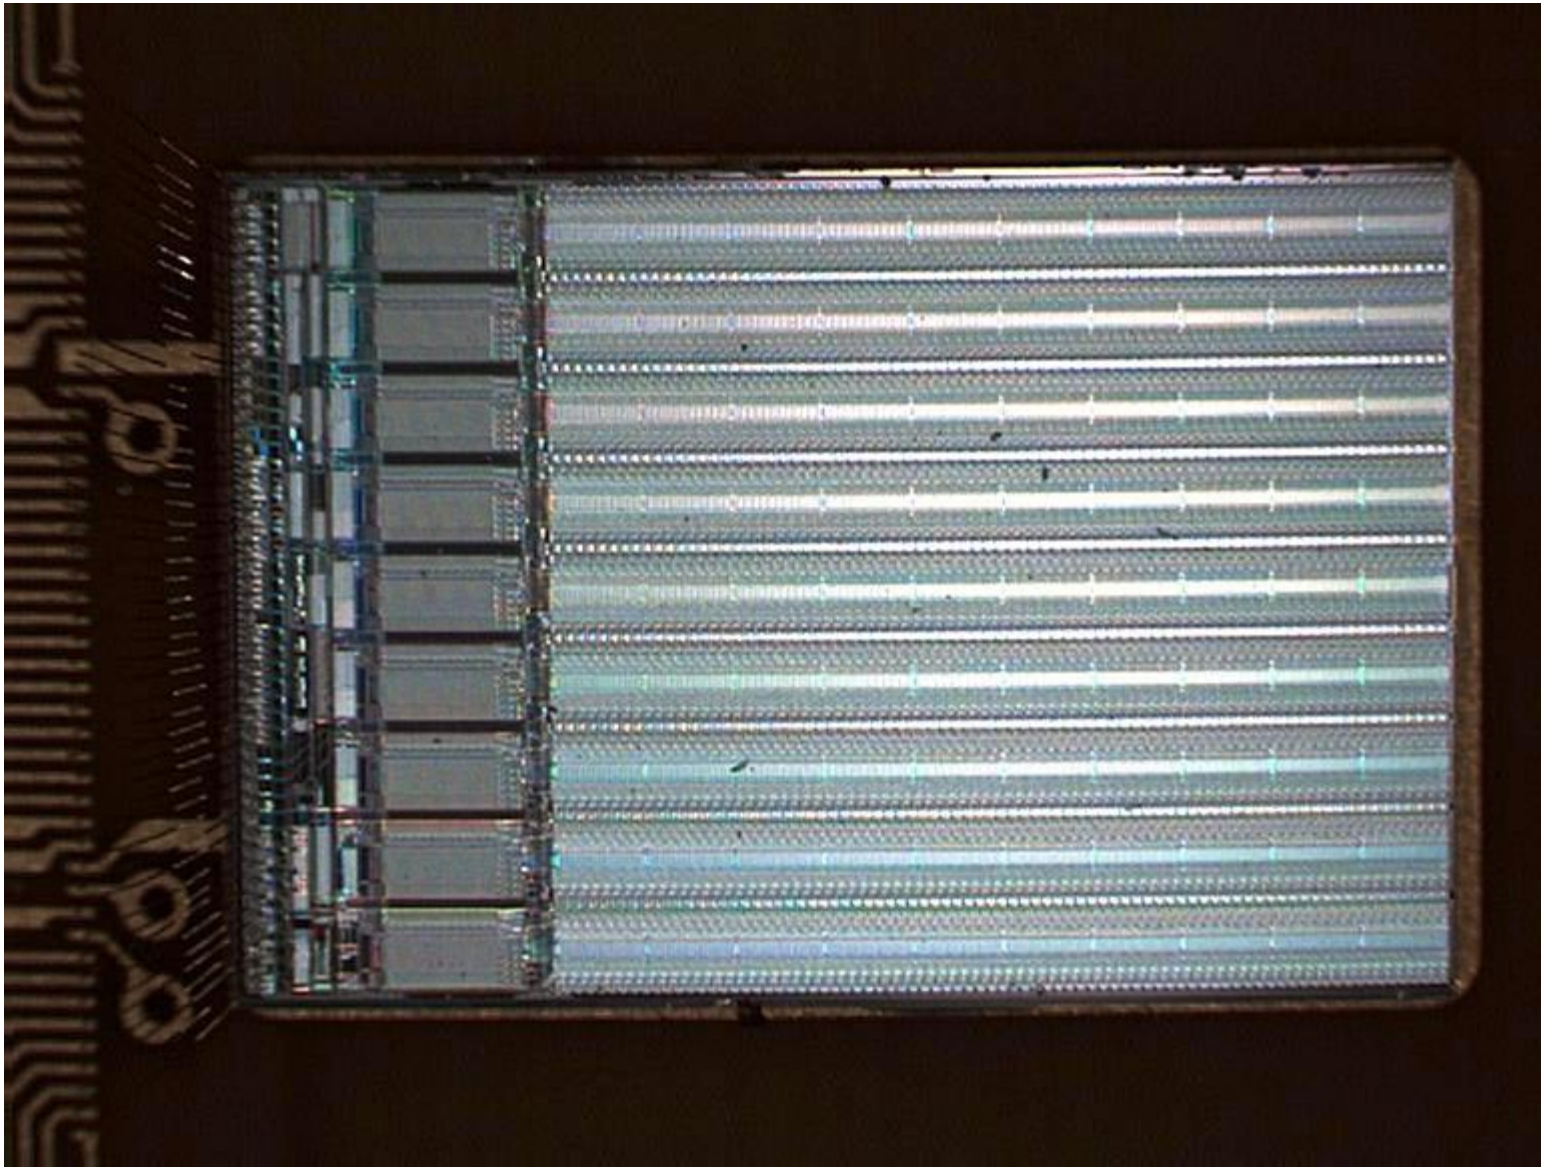

Expanded view of Bottom of FE-D1 chip:

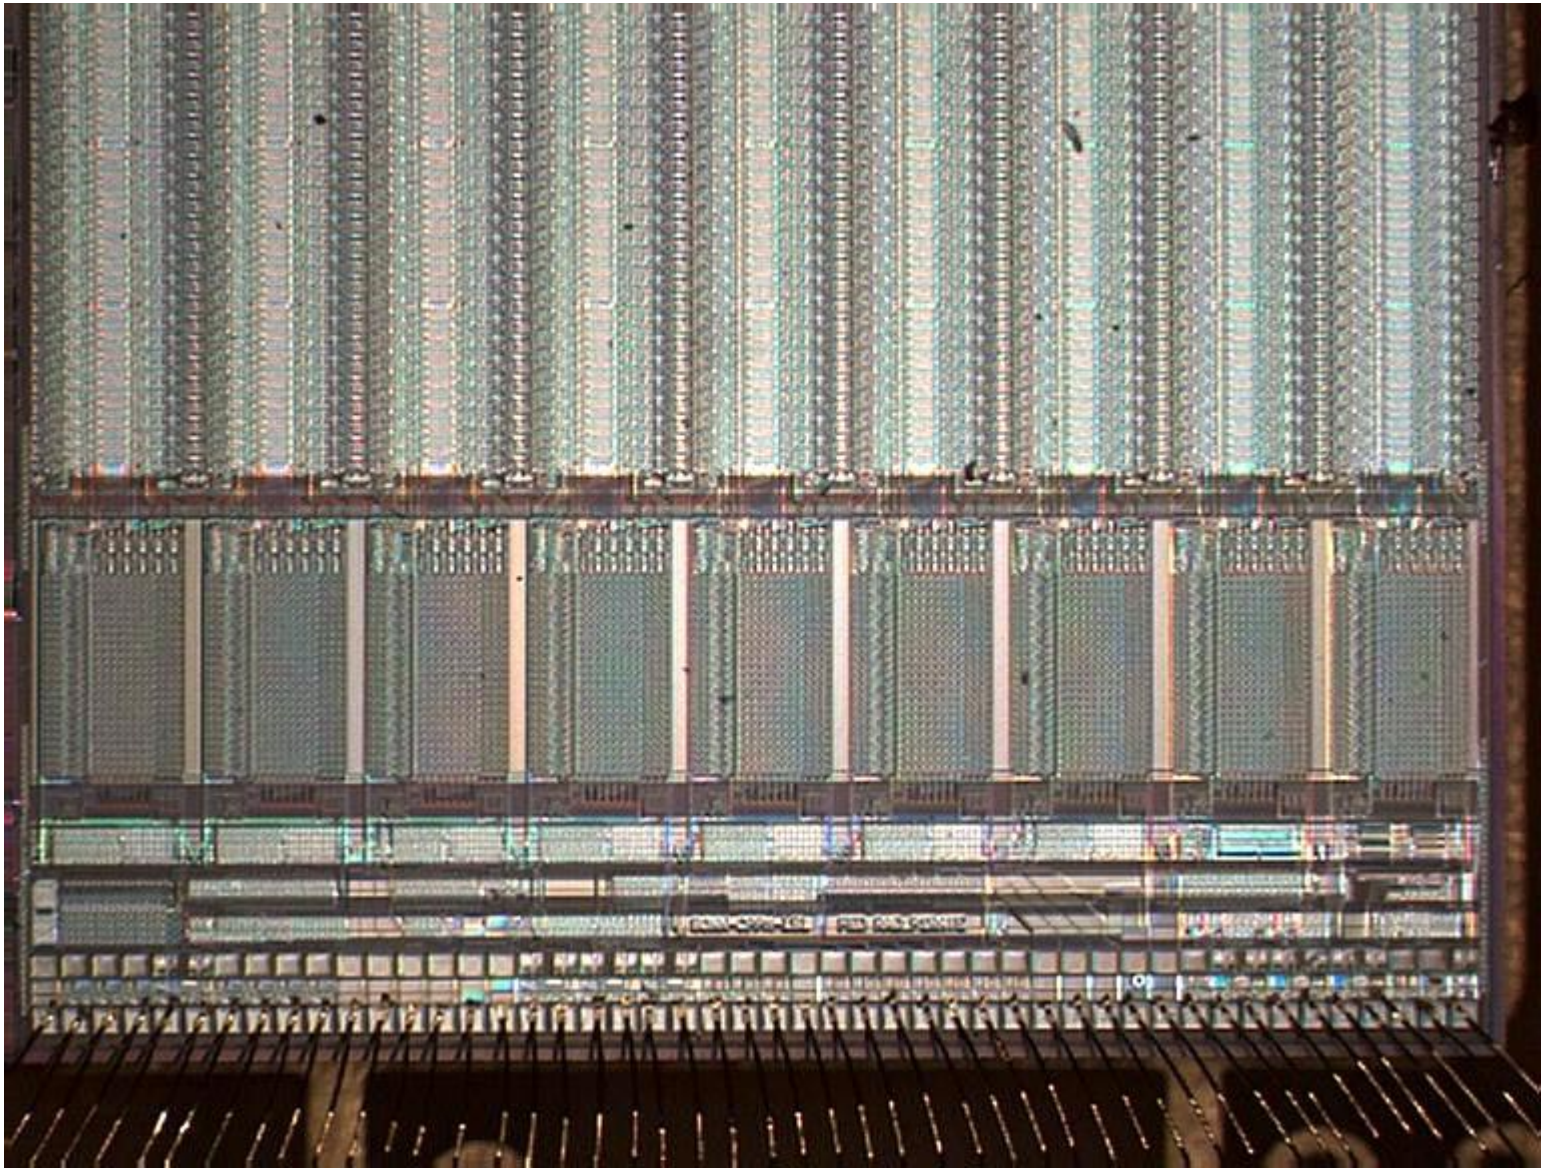

Front-end, biassing and control

Summary of the requirements:

- A nominal capacitive load of 200 fF is expected, roughly half to ground (parasitic) and half to the nearest neighbors (inter-pixel). Good performance should still be obtained with loads of 400-500fF. The n^+ on n-bulk detectors provide negative signals.
- Pixels are oriented to maximize signal and efficiency (minimize charge sharing).
- The outer layers should be 250μ silicon, and the B-layer should be 200μ.
- The expected signal after the lifetime dose of 10^{15} n-equiv/cm^2 is about 6Ke with 200V bias, and about 10Ke with 600V bias. We propose to operate at the higher bias at the end of the detector lifetime, and have real prototype experience to show that this works well.
- This leads to an in-time threshold requirement of about 4Ke (this would be 2.5Ke for the lower bias voltage). This requirement has often been defined using a maximum timewalk relative to some large reference charge of 25ns, however, we believe 20ns is the right target for the complete FE chip. This could be achieved by for example setting a 3Ke threshold, and having the required overdrive for a timewalk of 20ns be less than 1Ke. This is the most challenging requirement for our front-end.

- Noise should be less than 300e and threshold dispersion less than about 200e, leading to an overall threshold “variation” of less than 400e.
- Leakage current tolerance should be at least 50nA per pixel, without significant changes in operating performance, and independently achieved for each pixel.
- Noise occupancy should be less than 10^{-6} hits/crossing/pixel.
- Crosstalk between neighboring pixels should be less than 5-10%, where this is defined as the ratio between the threshold and the charge which must be injected into a pixel to fire its neighbors.
- A double pulse resolution of 2μs is required for the outer layers, and 0.5μs for the B-layer, in order to achieve our total deadtime requirements.
- It is required to provide binary readout of each pixel, but a modest analog resolution (4-5 bits) is very desirable if it can be achieved without a large impact on the other performance specifications.
- A threshold range of 0 - 6Ke is needed.
- A calibration injection capacitor of 10fF should be included in each pixel.
- We do not know whether real diode input protection is needed. In all present chips, no explicit input protection is provided, and no major problems have been observed.

FE and Control Blocks:

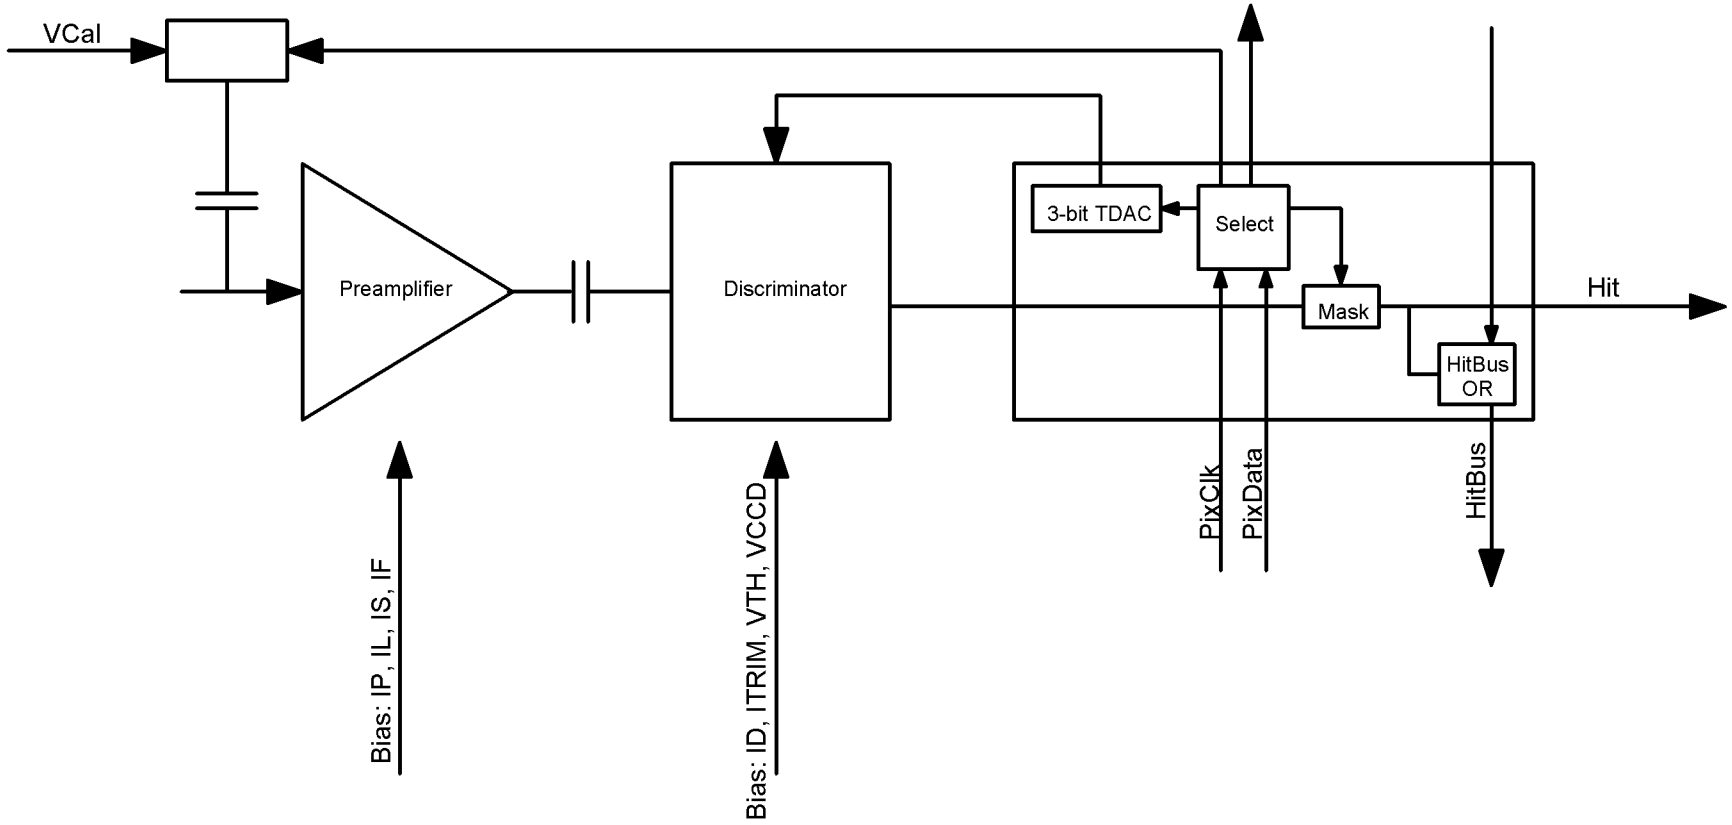

- Preamp has roughly 4fF DC feedback design, 12fF injection cap, and 40ns risetime. Input transistor operates at about 9µA bias and 1.5V. Cascode and follower operate at 3V and about 1.5µA each.
- Discriminator is AC-coupled and uses a bias current of about 6µA, plus two tunable voltage supplies (VTH and VCCD) which control the threshold and AC-coupling time constant.

FE Biassing and Control Blocks:

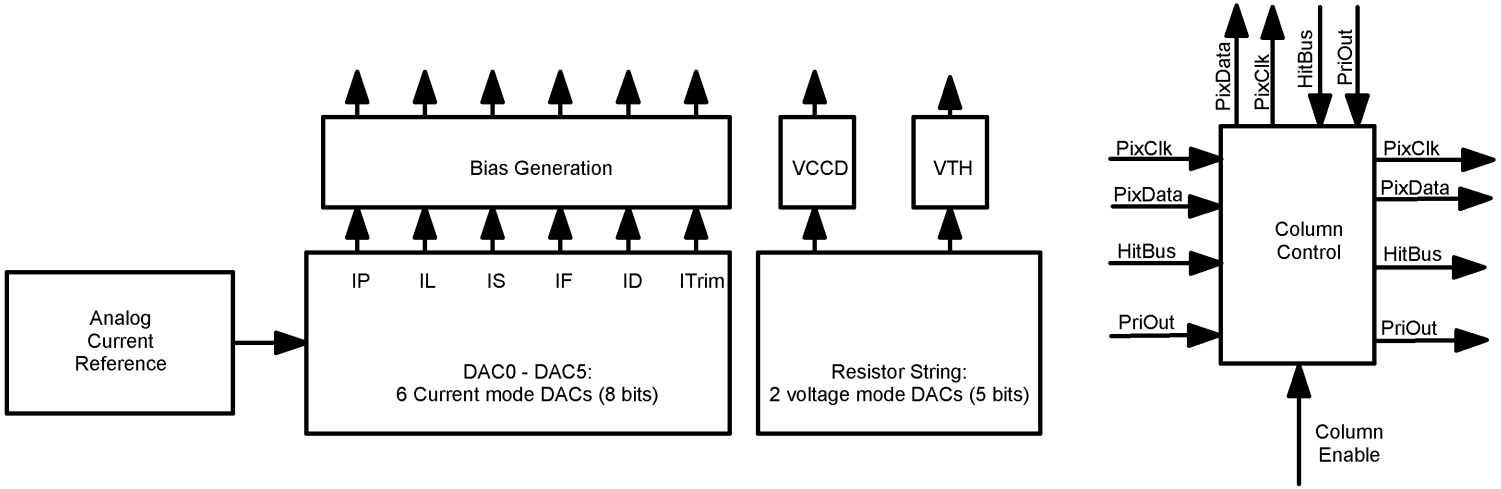

- A V_{BE} reference is used to supply a 4μA reference current to the current mode DACs (1μA/bit, but two LSB generated by simple mirrors).
- The current mode DACs are a rad-tolerant 8-bit design with good linearity.
- The two tunable voltage supplies use a common resistor chain to control output amplifiers which supply the necessary voltages with up to about 5mA of drive.
- A single column enable bit controls the major operations of a column pair. It allows bypassing a column pair in the pixel shift register chain, bypassing a column pair in the HitBus FastOR net, and bypassing the sparse scan readout of a column pair when transfering an event out of the FE chip.

Digital readout

Summary of the requirements:

- Make a unique association of each hit pixel with a 40 MHz beam crossing.
- Store hits in pixel array for L1 latency period, which can extend up to 3.2μs.
- Make a modest TOT measurement by counting time differences between leading and trailing edges in 40 MHz units.
- Simulations for the current architecture exist, driven by the full GEANT simulation of ATLAS. This suggests that the current architecture needs to operate with a 20 MHz column clock rate and have 25 buffers per column pair in order to provide safe operation of the outer layers. The B-layer requirements are more stringent, and require something like 32 buffers. Further study is underway for this layer.
- There is only a single error condition which occurs, namely overflow of the EOC buffers. In the case where the EOC buffer block in a given column pair overflows, hits are lost until a free buffer exists, and the error condition is stretched to cover a full L1 latency (covering all possible events which could have lost hits due to this condition). The error status is then transmitted in the EOE word.

Specifications:

- Clock duty cycle specified to be between 40% and 60% (high phase lies between 10ns and 15ns, or nominal +/- 2.5ns).

Block diagram of the basic column-pair readout:

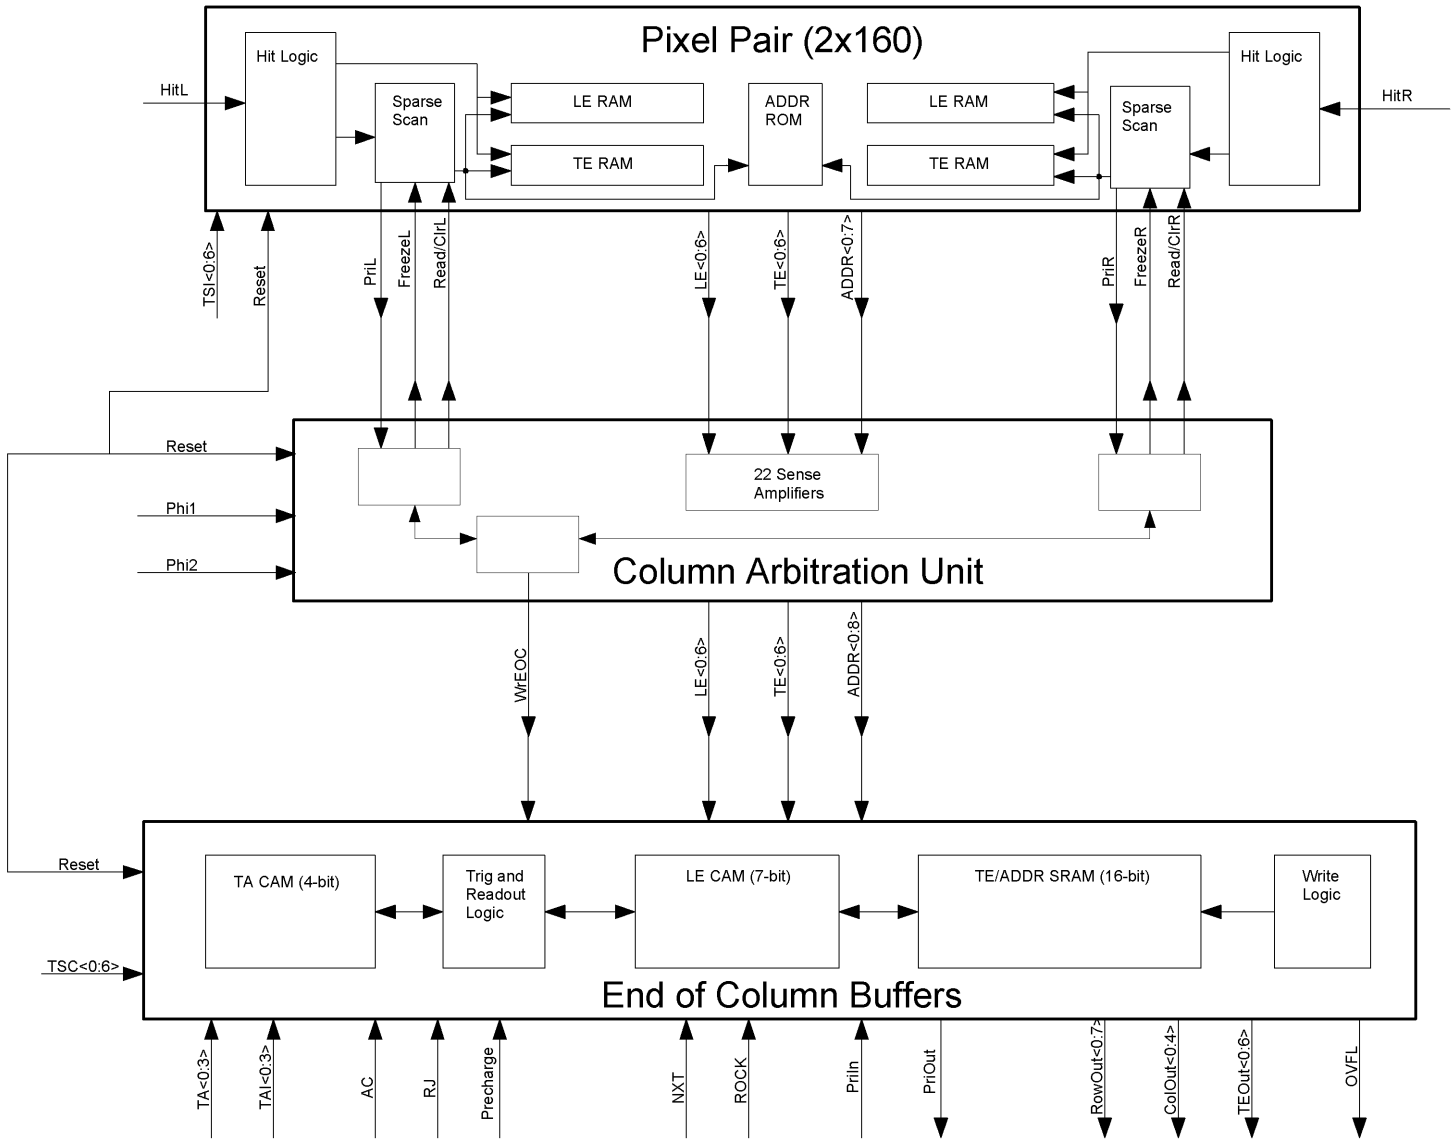

3

Summary of basic steps in readout of pixel data:

- Transfer hit information (LE and TE timestamp, plus pixel row address) into an EOC buffer. This operation begins when data is complete (after discriminator trailing edge). The transfer of hits from a column pair is synchronized by the CEU in the bottom of column, which operates at a speed of 5, 10, or 20MHz.
- Hit information waits in the EOC buffer for a corresponding LVL1 trigger. If a trigger arrives at the correct time (checked using LE timestamp of hit), the data is flagged as belonging to a particular 4-bit trigger number. Otherwise it is reset.
- Once the chip has received LVL1 triggers, the trigger FIFO will no longer be empty. This initiates a readout sequence in which the EOC buffers are scanned for the presence of hits belonging to a particular trigger number. If hits are found, they are transmitted to the serializer. After all hits are sent, an End-of-Event word is appended to the data stream.
- All of these operations occur concurrently and without deadtime.

Transfer of pixel data from pixel to EOC buffer

- The leading edge of the discriminator in a pixel causes the current timestamp to be latched into 7-bit RAM in pixel. The pixel is flagged as hit using an RS FF.
- The trailing edge of the discriminator causes the current timestamp to be latched in a second 7-bit RAM. This initiates the sparse scan operation.
- A scan ripples down the column, in 10 blocks of 16 pixels each. After the scan has had time to settle, the topmost pixel in the column will be selected for readout (and no other pixel will be selected).
- Meanwhile, the CEU should have seen the results of the sparse scan, indicating that there is a hit pixel in the column, awaiting readout. The CEU responds by issuing the Freeze signal to prevent any new hits from entering the sparse scan. The hit status inside pixel is propagated to dynamic node known as FrozenHit.
- The CEU issues the ReadPix signal to whichever of the two columns in the pair it has decided can transfer data. This causes the pixel to gate its LE/TE/Addr information onto a bus with senseamps at the bottom of column. It also releases the sparse scan to run in parallel with the data transfer and select the next pixel for readout, allowing hits to be scanned out at full column clock frequency. The CEU also issues the WrEOC signal, and data is stored in a selected free buffer.
- The CEU issues the ClearPix signal, reseting the hit FF of the selected pixel.
- This continues until the sparse scan indicates no pixels are left to be read out.

Operation of EOC Buffers

- Buffer has three status FF, one for “occupied”, one for “triggered”, and one for “reading out”. Initially, all buffers are free, with none of these FF set.
- One buffer consists of 16 bits of SRAM for TE/Addr data, a 7-bit CAM for LE data, and a 4-bit CAM for Trigger number, plus logic to support the simple state machine that governs the operation of each buffer. CAM blocks are write-only
- Once a buffer is selected to be the “next free” buffer, the next WEOC signal will store a hit in it, and it will be marked as “occupied”. If the buffer block is full, no further writes can occur, and an overflow signal is generated.
- This starts the operation of the 7-bit LE CAM block, which looks for a match with the trigger timestamp (displaced by L1 latency from timestamp seen by pixels). This block compares timestamps at 40MHz. If a match is detected, and there is no trigger present, the buffer is declared free. If there is a trigger present, the buffer is flagged as “triggered”.
- The readout phase is initiated by the presence of trigger information in the trigger FIFO. A scan is performed using the 4-bit CAM to find all of the hits in the EOC buffers which belong to the same trigger as the one being read out. If a match is detected, the buffer is placed in the “reading out” state.
- These hits are then read out using a 2D sparse scan to select one hit at a time. After hit readout, buffer status is cleared. After all hits are sent, EOE word is sent.

Block diagram of the Readout Control:

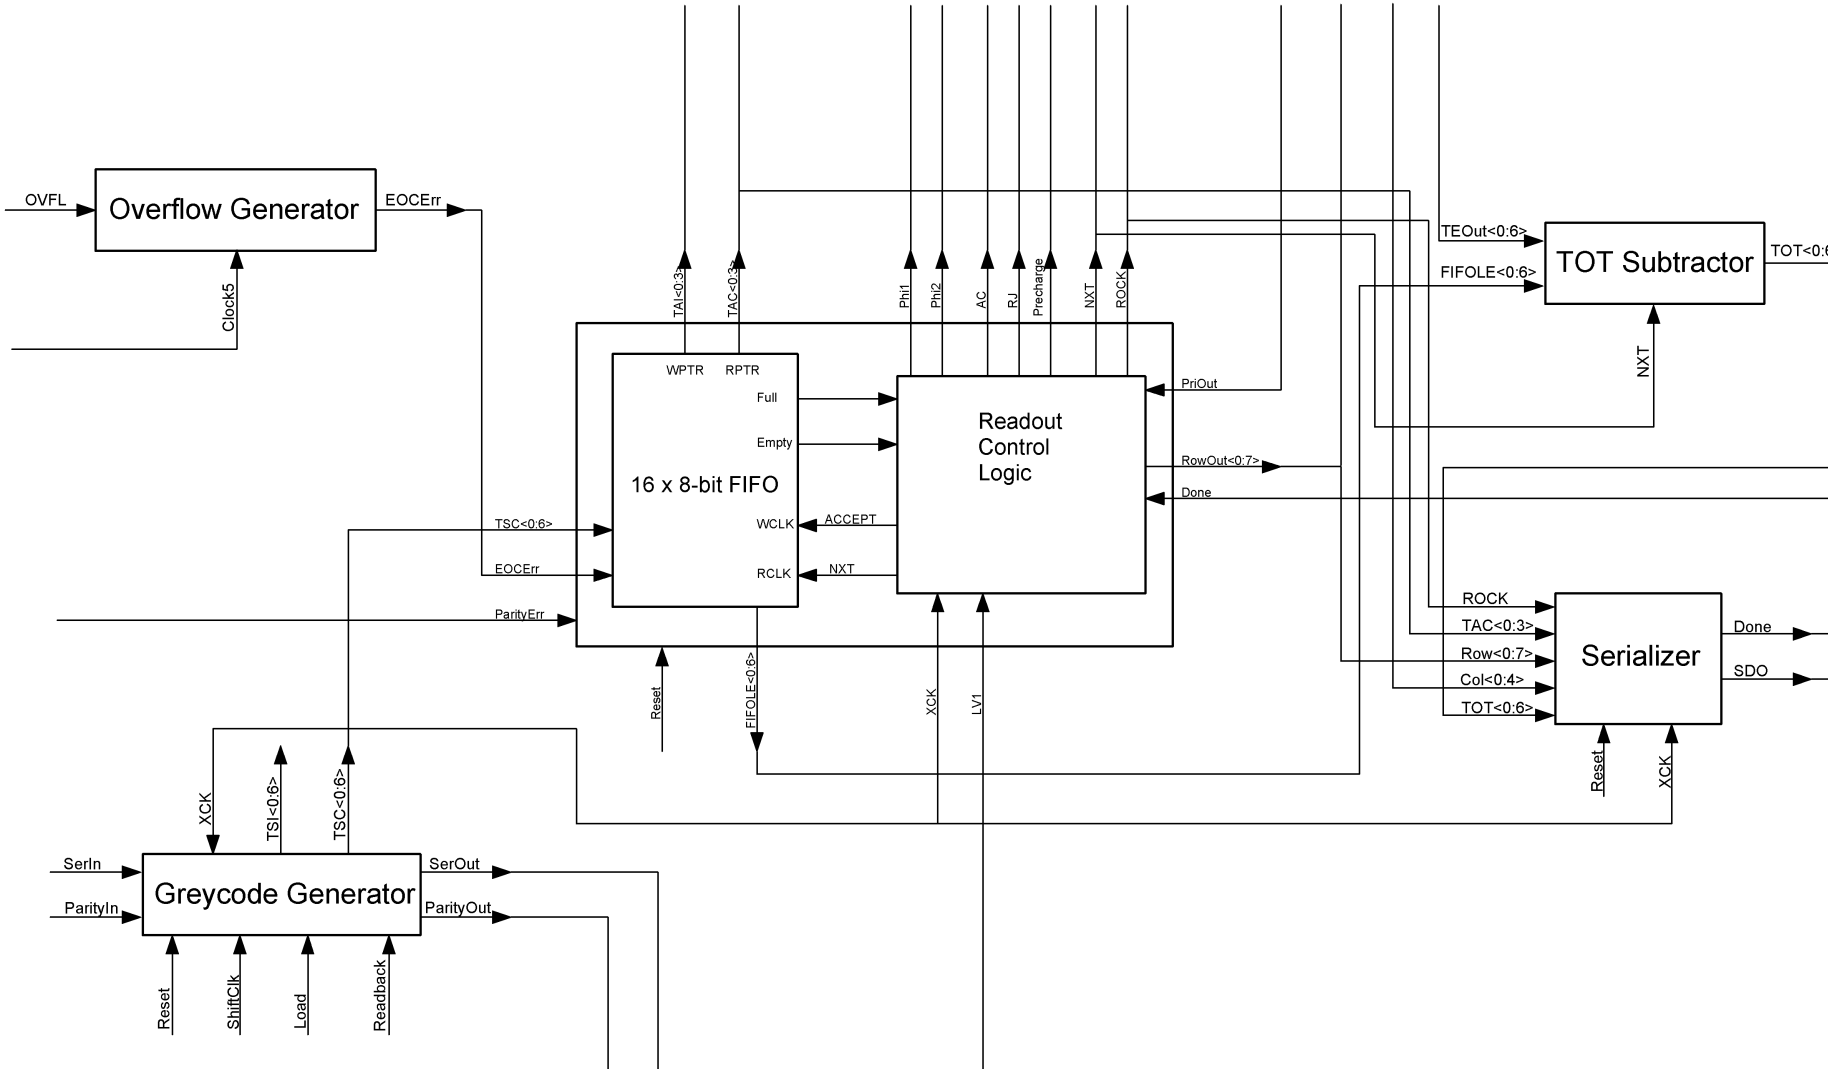

Block diagram of Command Decoder and Global Register:

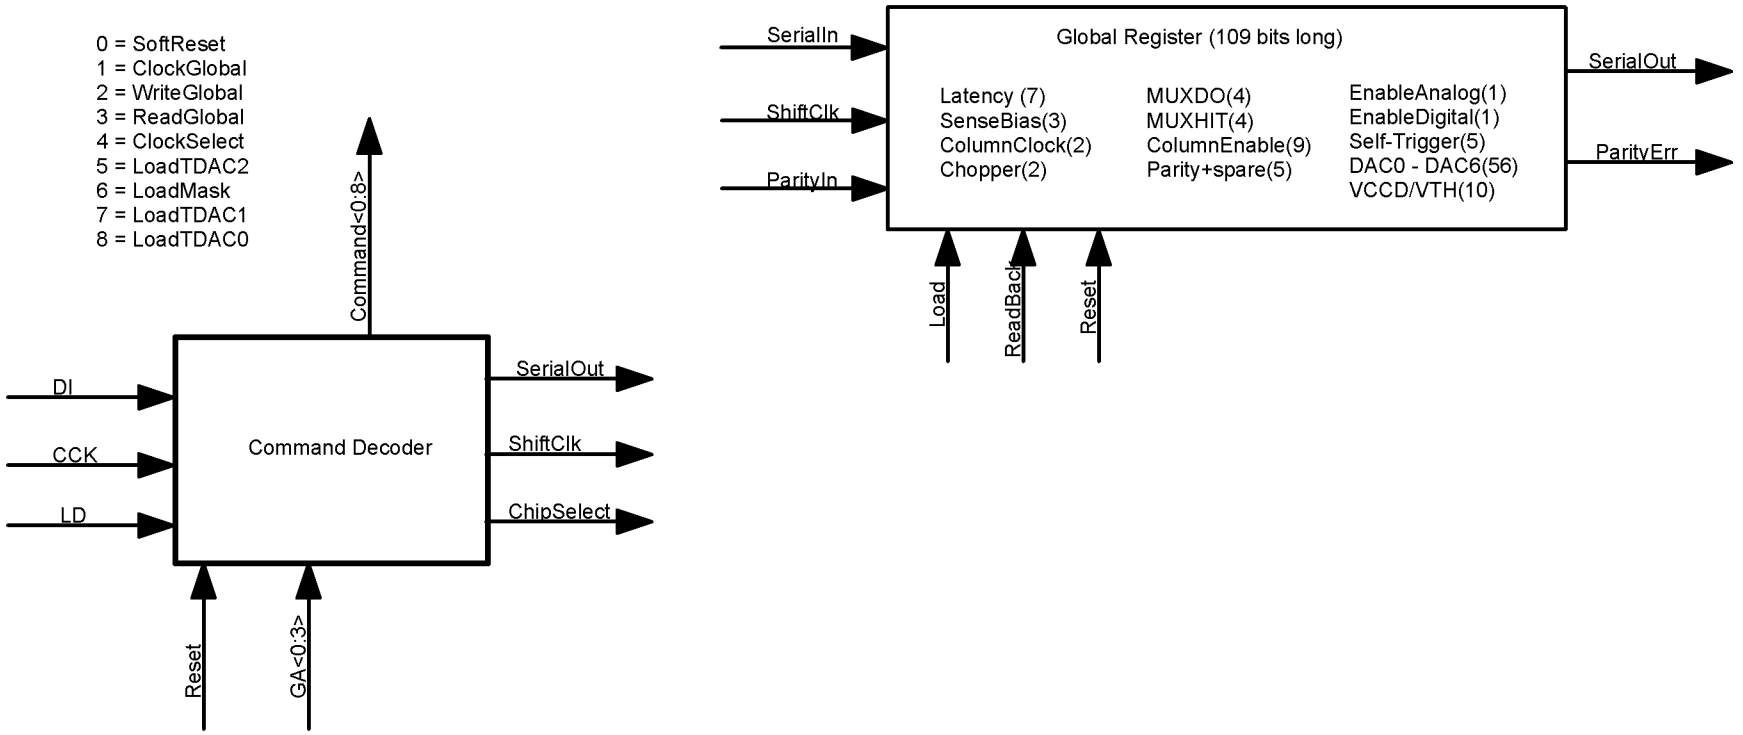

- Simple command protocol, based on a 5+9-bit command field, after which LD goes high and associated data may be transmitted. This supports 9 different, independent commands.
- Global Register controls overall operation of FE chip. Because of the critical importance of its bits, it does a parity calculation (chain of XORs) of the contents, and generates a Parity Error if the parity is not what is expected (due to SEU). This parity error is transmitted as part of the EOE word to flag such errors.

Block diagrams of remaining blocks in the periphery:

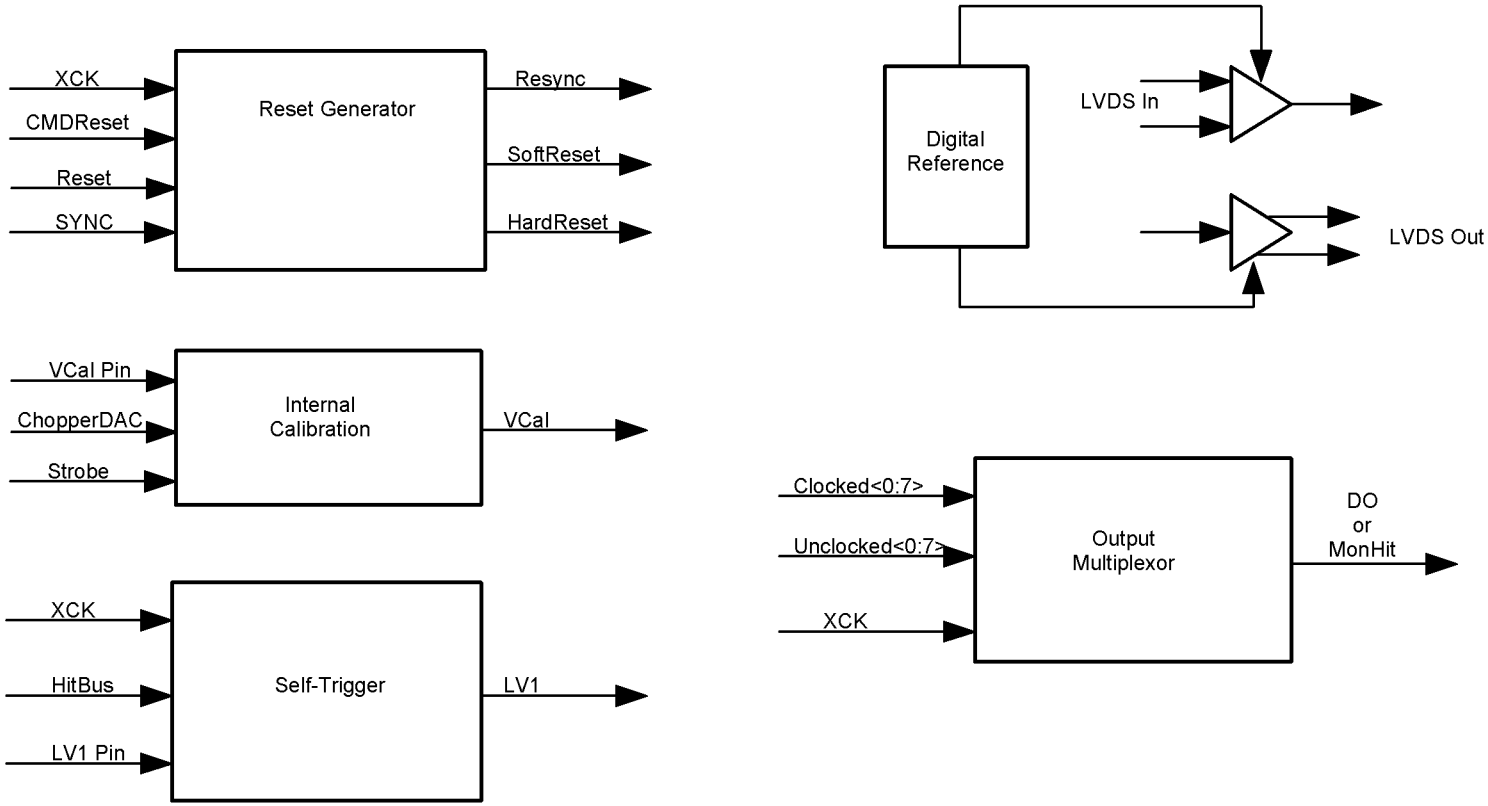

- FE-D uses a steered current of up to 5mA, applied across a precise resistor, in order to generate the voltage step for the injection capacitors. This provides a better pulse shape than a large FET voltage chopper, and allows internal generation of the current using our standard 8-bit current DAC.
- There is a reset generator that either uses the external RESET pin, or the width of the SYNC input in XCK clocks, to generate internal reset signals.

- There are three basic reset signals. The Resync makes sure that all FE on a module are using the same trigger number (it resets the trigger FIFO). The SoftReset puts the chip into the “empty” state for data, but does not alter any configuration information. Finally, the HardReset also resets all configuration information to zero.
- There is a self-trigger generator, which either passes the input LV1 through to the trigger processing circuits, or uses the internal FastOR signal to generate its own LV1 signal after a programmable latency. This allows the chip to be used in self-trigger mode with a source, and it will produce output data once it has been armed by a previous LVL1 trigger, and it sees a signal on the internal FastOR.
- There is a dual 8-fold output multiplexor which selects which internal signal or data stream is transmitted off the chip through the serial output. The first eight inputs are synchronized with XCK, while the second eight are not. An identical multiplexor circuit is used both for the standard serial output pins, and for the MONHIT output pins.
- The LVDS driver/receiver circuits use a second internal current reference to define the drive current. The common mode voltage is referenced to the DVdd supply using resistors. While not quite conforming to the LVDS specification, the design works well for local pixel module communication. Have performed corner simulations of the performance of a driver/receiver pair with realistic loading. The most critical signal (XCK) was found to be within specifications.

Summary of Problems with FE-D1:

Layout Problems:

- One of the buffer amplifiers used in the voltage DACs for front-end control had a layout problem. This should have been caught by the LVS/DRC checking, but was not. The result was that a short was created between two capacitors in the VTH amplifier, and it did not work properly. For testing, the amplifier output was isolated by FIB modification, allowing the voltage to be driven externally.

Buffering Problems:

- Large buffers were missing from many of the outputs from the command decoder, resulting in very poor risetimes on these control lines. For chip testing, this was compensated by artifically slowing down the command operation.
- For the long Pixel Register, a single clock buffer for the two-phase clock was used (nominal frequency of 5MHz). The sizing of this buffer was a bit small for the actual capacitive load (estimated at almost 200pF). However, the resistance of the traces was neglected in the simulations. The risetime of the clock is poor (about 30ns), due to the combination of the large capacitive load and the resistive traces. This can only be fixed by distributing the buffers (large buffer at bottom of chip, 18 distributed buffers at bottom of column). There appear to be no operational effects of this for un-irradiated die, but if the risetime degraded further, the shift register would most likely fail.

- For the internal XCK distribution, there was a missing buffer after the LVDS receiver, so the receiver had to drive a roughly 4pF load, and the risetime was poor (about 5ns). This appears to have no operational effects (see later discussions of simulations).
- For the long connection of the 40MHz serial data stream from the Serializer output to the Output Mux, there was a missing buffer. The result was that an un-buffered FF was driving a load of about 1pF. In this case, if VDD was too low (below about 4.5V), this FF did not have enough drive to latch ones, and so bits would get lost. Extensive simulation, and operation of FIB-modified chips in the lab, confirmed that once this problem was bypassed, the rest of the digital readout circuitry operated properly at the nominal VDD of 3.0V.

Minor design flaws uncovered by corner simulations:

- Small glitch in CEU for *iss* simulation caused by excess loading on one node. Resulted in a runt WEOC pulse which wasted buffers in EOC. Fixed in FE-D2.
- Not well-optimized sparse scan device sizing, resulting in slightly slower sparse scans in both the column pair and the EOC buffers. Only an issue for *iss* case, and slightly modified only in column pair of FE-D2S.
- Undersized driver in EOC logic results in marginal operation of 7-bit LE CAM for *iss* case. This is presently where the digital readout breaks first in simulation. Not fixed in FE-D2 due to lack of space.

Design Errors:

- Forgot to mask overflows for disabled column pairs, so if a column pair is disabled, chip usually shows buffer overflow status for every event.
- Initial concept of self-trigger did not fit within MCC and test system requirements. Need to add concept of “arming” of self-trigger, as the rest of the system cannot deal with a completely autonomous data source.

Analog problems:

- There are several strange problems with the analog performance of complete chips (not observed in the analog test chip).
- The first is anomalous noise behavior observed, particularly for bump-bonded assemblies. The noise value is larger than expected, and has a strong dependence on preamp feedback current. Some of this can be attributed to digital cross-talk between the readout logic and the sensor, coupling back in to the preamplifier input. The measurements are confusing, not yet conclusive.
- A second problem is an instability in the threshold values (shifting by few hundred electronics, and systematic geographic variations which themselves vary, when hard resets are sent to the chip under test. No mechanism found for this.
- A third problem is an apparent systematic geographical variation in the TOT value observed for a fixed input charge. This does not seem to be correlated with any known bias distribution problem in the layout.

Other Problems (to be discussed in detail later):

- There are several serious problems that don't fit into the above categories. They appear to arise from a high rate of defective devices, when these devices are used in dynamic logic blocks.
- The first problem is the very poor yield observed for the Pixel Register (2880-bit two-phase, quasi-static control register which snakes through the active part of the pixel array). The yield for wafer-probed FE-D1 die, for this test alone, was about 25%. For FE-D1b die, this improved to the range of 80-90%, but is still very poor for circuitry which represents only a few percent of the die area.
- The second problem is the significant disruption of the digital readout of a complete column pair by a certain class of defective pixel. After detailed study, we concluded that this problem was largely related to a high defect rate in a particular NMOS transistor used in a dynamic node in the pixel hit logic.

Studies of Pixel Register problem:

- Studies of the pattern of bit errors that occur indicate that the defect is most likely related to anomalous leakage from the dynamic node in the FF, and most likely leakage across a particular PMOS transistor.
- Studies have been performed at reduced shift register clocking frequency, and additional problems do appear. These would be consistent with a spectrum of leakage defects.
- Detailed study of this problem is difficult because the minimum segment of the register that can be studied is 320 bits long. There is no simple way to localize a defect below this level, or to see more than one defect in this register segment.

Studies of defective pixels in the column pair readout:

- For this problem, it is possible to localize the defective pixels. In many cases, providing a hard reset to the digital circuitry will clear the problem, but some very defective pixels cause problems after reset even when their input is disabled.
- The symptoms of the problem are complex, but allow us to constrain the possible faults that could cause them. It appears that the problem arises when the NMOS FET used to reset the FrozenHit node in the pixel hit logic is leaky (that is, has an off-resistance of less than 1MΩ). Probing of gate and drain nodes of this device, using FIB-deposited pads, confirms the electrical model we initially proposed, with bad pixels having NMOS with off-resistances of about 300KΩ.

Summary

- We have performed detailed studies of FE-D1 at wafer and die level, including bump-bonded assemblies with detectors. This has allowed us to analyze obvious errors in FE-D1 chips. These errors have all been corrected for FE-D2 designs.
- There are additional problems that we believe are related to the technology. We have been working with TEMIC to understand these faults, but with very limited success so far.

Some actions taken so far:

- Leakage measurements were made on all of the test structures on several wafers from the FE-D1 run by TEMIC. Leakage currents were all sub-nA, with no indication of the types of defects we see in our chips.
- One chip which had been fully characterized for defects at LBL was sent to TEMIC, who reverse-engineered the chip in the region of a defective pixel. This analysis indicated no particular problems with the poly etching, neither inside the active area (gates), nor in the routing nearby (shorts of the type observed by ATLAS LArg team in their SCA chip).
- We have been discussing with one of the process designers (R. Truche from LETI/CEA) for further suggestions, and he has also raised issues of antenna effects causing problems with the gate oxide, but in this region of FE-D, the antenna ratios seem to be too small for this (50 or so).